Uncovering allosteric pathways in caspase-1 with Markov transient analysis and multiscale community detection†

B Amor*[a,b], **S N Yaliraki**[a,b], **R Woscholski**[a,b] **and M Barahona***[a,c]

Allosteric regulation at distant sites is central to many cellular processes. In particular, allosteric sites in proteins are a major target to increase the range and selectivity of new drugs, and there is a need for methods capable of identifying intra-molecular signalling pathways leading to allosteric effects. Here, we use an atomistic graph-theoretical approach that exploits Markov transients to extract such pathways and exemplify our results in an important allosteric protein, caspase-1. Firstly, we use Markov Stability community detection to perform a multiscale analysis of the structure of caspase-1 which reveals that the active conformation has a weaker, less compartmentalised large-scale structure as compared to the inactive conformation, resulting in greater intra-protein coherence and signal propagation. We also carry out a full computational point mutagenesis and identify that only a few residues are critical to such structural coherence. Secondly, we characterise explicitly the transients of random walks originating at the active site and predict the location of a known allosteric site in this protein quantifying the contribution of individual bonds to the communication pathway between the active and allosteric sites. Several of the bonds we find have been shown experimentally to be functionally critical, but we also predict a number of as yet unidentified bonds which may contribute to the pathway. Our approach offers a computationally inexpensive method for the identification of allosteric sites and communication pathways in proteins using a fully atomistic description.

Keywords: allostery, allosteric pathways, community detection, complex networks, random walk, multiscale, caspase-1.

arXiv:1411.2847v1 [q-bio.BM] 11 Nov 2014

1 Introduction

Allostery describes the widely observed phenomenon by which a perturbation at one site of a protein has a functional effect at another, distant site[1]. Traditionally, studies of allostery have been linked to the cooperativity observed in large multimeric proteins such as haemoglobin. In this context, the classic 'induced-fit' (Koshland-Nemethy-Filmer, KNF)[2] and 'pre-existing equilibrium' (Monod-Wyman-Changeaux, MWC)[3] models both consider that each monomer has a high-affinity ('relaxed') R-state and a lower affinity ('tense') T-state. In the induced-fit model, binding of one subunit *drives* the next subunit into its new higher affinity R-conformation. In the MWC model, the protein ensemble is already in equilibrium between the T and R states, and binding of the ligand to one subunit *shifts* the equilibrium of the ensemble towards the high-affinity state.

The so-called 'new' view of allostery, which is essentially an extension of the MWC model to general allosteric effects, regards the allosteric effector as shifting the equilibrium of a pre-existing ensemble of conformations towards the less populated state[4]. In this sense, any protein could be allosteric[5], and the perturbation could be anything that changes the free-energy landscape of the protein (including perturbations which do not induce a visible conformational change[6,7]). This perspective views proteins as highly dynamic, with the ability to sample their active (i.e., less populated) state even in the absence of a ligand or substrate[8].

These thermodynamic models provide a helpful phenomenological description but do not explain how the change in the energy landscape is induced by the effector, and how this effect physically propagates between the effector site and the active site[1,9]. The idea of allosteric pathways (i.e., sets of contiguous residues through which a signal propagates[10]) has grown in popularity since Lockless and Ranganathan identified a set of evolutionarily conserved residues linking the binding site and a distal site in the PDZ family of proteins[11]. NMR[12,13] and molecular dynamics[14,15] studies on members of this family have suggested other overlapping pathways of energetically linked residues connecting the binding site to further distal sites. Further computational studies have used combinatorial unfolding of protein structures combined with free energy calculations to find distal regions of proteins which are 'energetically coupled'[16,17], but do not give a structural interpretation of this coupling.

Recently, residue-residue interaction networks (RRINs) have been used to model allosteric communication. Using such coarse-grained representations, Del Sol *et al* identified central residues that contribute most to reducing the average shortest path in the RRIN network[18], and showed that central

† Electronic Supplementary Information (ESI) available: See DOI: 10.1039/c4mb00088a

[a] *Institute of Chemical Biology, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK.*
[b] *Department of Chemistry, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK.*
[c] *Department of Mathematics, Imperial College London, South Kensington Campus, London, SW7 2AZ, UK.*
* *email: b.amor11@imperial.ac.uk, m.barahona@imperial.ac.uk*

residues often lie at the interface of modules in the RRIN[19]. These studies used classical extremal network measures such as betweenness to identify central residues, i.e., these measures are based on shortest path calculations. However, allosteric communication is characterised by multiple major and minor pathways[10]. Other studies have used *ad hoc* deterministic approaches to consider these, such as including pathways within some small distance of the shortest path[20,21]. Taking a stochastic perspective, Chennubhotla and Bahar[22] (followed by others[23,24]) used hitting and commute times of random walks on RRINs to explore intra-protein communication and linked these measures to the equilibrium fluctuations of a coarse-grained, residue-based Gaussian Network Model of the protein[22]. Stochastic methods[14,25] have also been used to analyse clustering in protein structures.

In this paper, we study the pathways implicated in allosteric regulation by considering the transients of random walks taking place on an atomistic graph representation of the protein. Our approach differs from the above methods in two crucial ways. Firstly, we use an atomistic rather than a residue level description, i.e., each node in our graph corresponds to an atom (rather than a residue) with edge weights corresponding to the actual strengths of the bonds between atoms. This allows us to quantify the contribution of specific atomic interactions to communication pathways. Secondly, we use a dynamics-based *multiscale* method, Markov Stability[26], to analyse the community structure in the protein graph across all scales in one sway, from chemical groups to protein subunits. In contrast, other community detection methods (such as Modularity[19,27]) find just one partition at a particular scale[28], and suffer from a resolution limit[29] that prevents them from identifying the multilayered community structure that exists at different scales in highly organised networks such as proteins. Our method overcomes this limitation because it scans across *all* scales identifying the levels of resolution where there is strong community structure. This is equivalent to the observation of a Markov transient over different time scales. Such transient behaviour can be explicitly used to explore the communication between specific sites in the graph. To do so, we introduce a measure for intra-protein communication, the characteristic transient time $t_{1/2}$, which takes into account *all* possible pathways between the communicating sites. We then use $t_{1/2}$ to identify groups of atoms that are strongly linked to the active site, as well as bonds that are key participants in the communication paths. In contrast to shortest path methods, our approach considers the contribution of all possible pathways between the two sites. Furthermore, the computational efficiency of our method enables us to carry out full mutational analyses to evaluate the relevance of all bonds and residues in the structure.

As a case study, we consider here the cysteine protease *caspase-1*, an important enzyme in which extensive experiments have identified an allosteric site[30] and details of a communication pathway between the allosteric and active sites[31]. Caspase-1 processes the pro-inflammatory cytokine interleukin-1β and has specificity for substrates with aspartic acid adjacent to the peptide bond being broken[32]. Members of the caspase family are involved in signalling pathways associated with apoptosis and inflammatory response, and as such are promising drug targets[33].

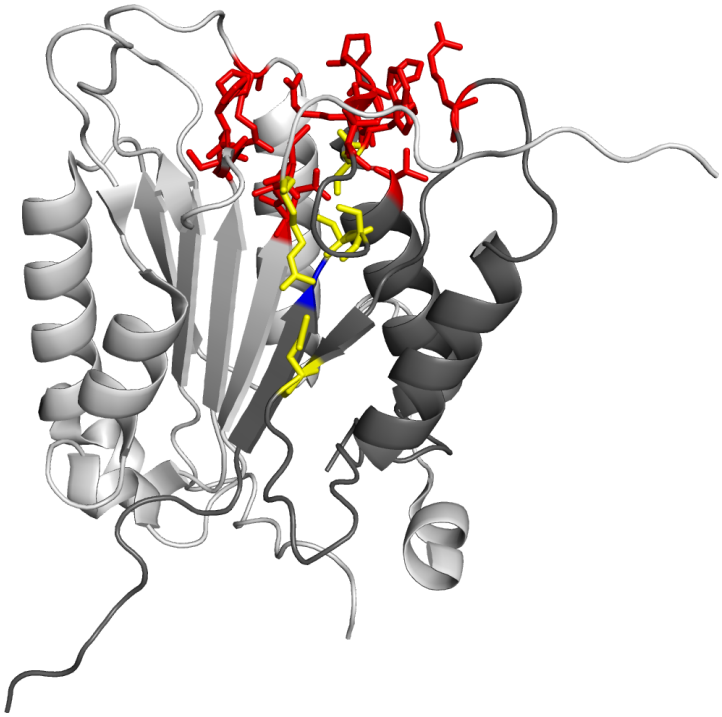

Fig. 1 Three-dimensional structure of the active conformation of caspase-1 (PDB: 2HBQ) showing the active site residues (in red), the allosteric binding site residue Cys331 (in blue), and the residues involved in the hydrogen bonding network (yellow sticks)[31]. The structure is a dimer formed by the large p20 subunit (light grey) and the smaller p10 subunit (dark grey). All protein figures in this paper were created with PyMol (http://www.pymol.org).

Caspase-1 is a dimer composed of a smaller p10 subunit and a larger p20 subunit[34], with experimental indications that two such dimers might combine to form a $(p20)_2/(p10)_2$ tetramer. Here we consider the dimer structure, which has been well characterised structurally. The active site spans across both subunits: residue Cys285 in the p20 subunit is the active site nucleophile, while the rim of the binding pocket is composed of residues 283-285 and 236-238 in the larger p20 subunit and residues 338-343 in the smaller p10 subunit. Scheer *et al* have identified experimentally an allosteric binding pocket situated at the dimer-dimer interface[30]. Datta *et al*[31] used structural data to identify a network of hydrogen bonds which link the allosteric and active sites in the active conformation (Fig. 1),

but are absent in the inactive conformation. A subset of those residues were then found to have a significant impact on catalytic activity. In particular, binding of the allosteric ligand disrupts a salt-bridge between residues Arg286 and Glu390, and an alanine mutation of either residue greatly reduces the catalytic activity of the protein.

Our analysis of the active and inactive conformations of caspase-1 shows that the active conformation has a less compartmentalised community structure than the inactive structure, which leads to increased communication between the active site and the rest of the protein. Through computational exploration of all point mutations, we find that only a few residues have a strong impact on this increased communication. We then consider the explicit analysis of transient random walks in the structures and find that the region of the protein with the largest increased connectivity in the active conformation corresponds to the allosteric site. Our computational mutational analysis then shows that the bonds that contribute most to this increased connectivity are formed by functionally important residues. Our method thus reveals the location of the allosteric site and the bonds involved in signal transmission directly from structural data.

2 Materials and methods

2.1 Structural data

We analyse three crystal structures of human caspase-1: one in unliganded ('inactive') form, and two complexed ('active') structures in complex with tetrapeptide substrates at the active site. The unliganded structure (PDB ID code: 1SC1) was obtained by Romanowski *et al*[35] through X-ray crystallography at a resolution of 2.6$\AA$. The complexed caspase-1/z-VAD-FMK (PDB ID code: 2HBQ) was obtained by Scheer *et al*[30] at a resolution of 1.8$\AA$. The structure of caspase-1 in complex with a tetrapeptide aldehyde inhibitor (PDB ID code: 1ICE) was obtained by Wilson *et al*[34] at a resolution of 2.6$\AA$. We follow the standard residue numbering as in Ref. Wilson *et al.* [34].

2.2 Construction of the atomistic network

In contrast to most network methods for protein analysis [18,19,22–25], our method starts by the construction of a fully atomistic graph representation of the protein. The graph is built from the structural information contained in the PDB file[36], which contains the Cartesian coordinates of each atom in the protein. Each node in the graph corresponds to a single atom and each edge defines a covalent bond or weak interaction (hydrogen bonds, salt bridges, or hydrophobic tethers)[37]. Any missing hydrogen atoms are added using the software Reduce[38]. We identify the presence of covalent bonds and weak interactions using the program FIRST[39] with a cutoff of 0.01 kcal/mol for hydrogen bonds and 8Å for hydrophobic interactions. Each edge has a weight which is linearly related to the bond energy. The bond strengths are obtained from the DRIEDING force-field[40]. This protein graph is encoded by a weighted adjacency matrix A, an $N \times N$ symmetric matrix (where N is the number of atoms in the protein) in which the entry A_{ij} gives the energy of the interaction between atoms i and j (0 if there is no interaction). For further details of the network construction, including the force fields used, see the Supplementary Information and Ref. Delmotte *et al.* [37].

2.3 Multiscale community detection with Markov Stability for protein structures

Intuition for Markov Stability analysis of atomic protein graph. Proteins are multiscale biomolecular machines with structural organisation at scales ranging from chemical groups, through amino acids, to protein domains, and even different subunits. This multiscale structural organisation is encoded in the atomic protein graph described above, which contains detailed physico-chemical and geometric properties of the protein. To reveal the organisation of the structure at different scales, we analyse the generated protein network using recently developed graph-theoretical techniques.

Markov Stability is a general method for multiscale community detection in graphs[26], and is thus well suited for the analysis of the protein graph at all scales. In contrast to other community detection methods (such as Modularity[27]), Markov Stability finds partitions of a graph into non-overlapping subgraphs ('communities') without imposing a particular scale *a priori*. Rather than obtaining a single partition, Markov Stability finds an optimised partition at every scale and decides if such a partition is significantly robust by establishing if the communities correspond to subgraphs where a random walk is likely to remain trapped over a certain timescale. As we increase this Markov timescale, the method acts as a zooming lens, scanning across all scales looking for significant communities at different resolutions. Hence, as the Markov time progresses, the partitions become coarser. In the case of the protein graph, this process allows us to scan across resolutions and find communities involving a few atoms (corresponding to chemical groups) at very short Markov times, through biochemical units (amino acids) and secondary structures (helical turns) at intermediate Markov times, to communities corresponding to conformational groupings or protein subunits at long Markov times. For a detailed description of the method, see the Supplementary Information and Refs. Delvenne *et al.* [26], Schaub *et al.* [29], Delmotte *et al.* [37].

Optimal partition of a graph into communities at a given scale. We now make these notions more precise. A random walk on a graph is described by a $N \times N$ Markov transition

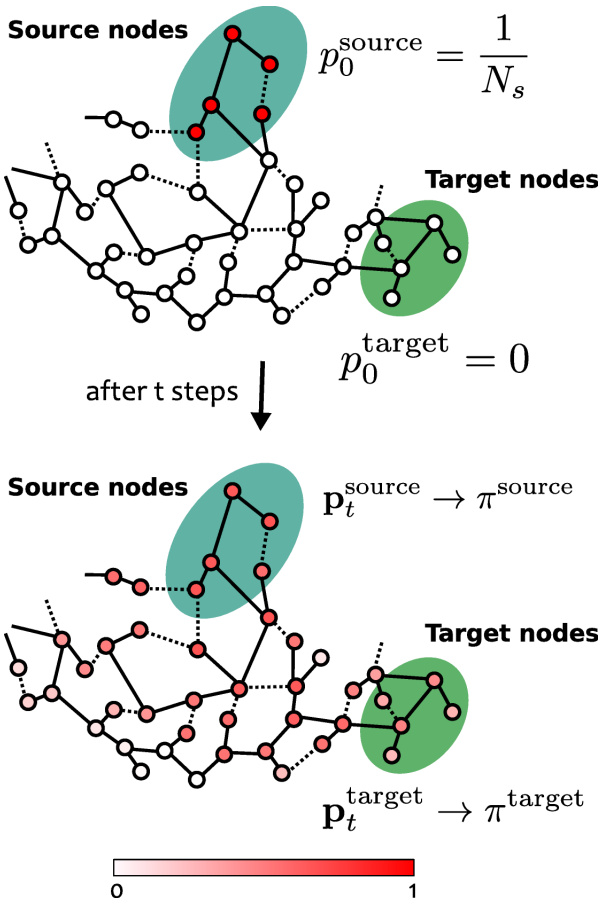

Fig. 2 Schematic illustration of random walk analysis. Initially the probability is distributed only across the source nodes. Through the iteration of the Markov process (1), the probability spreads across the other nodes in the protein. We monitor the time evolution of the probability at a set of target nodes and measure how long it takes for this set to reach half their stationary value.

matrix M, where N is the number of nodes in the graph:

$$\mathbf{p}_{t+1} = \mathbf{p}_t M. \tag{1}$$

Each element M_{ij} gives the probability of transitioning from node i to node j in one time step and $\mathbf{p}_t = (p_t^{(1)}, \ldots, p_t^{(N)})$ is a $1 \times N$ probability vector recording the probability of the process at each node at time t. M is directly related to the adjacency matrix by $M = D^{-1}A$, where D is the diagonal matrix of node degrees, i.e., D_{ii} is the sum of the weights of all edges incident to node i.

The key matrix for Markov stability is the block 'autocovariance' matrix [26]

$$R(t,H) = H^T(\Pi M^t - \pi^T\pi)H. \tag{2}$$

R is a $c \times c$ matrix where c is the number of communities in the partition, and $[R(t)]_{ij}$ gives the probability of the random walker starting in community i and finishing in community j after t timesteps. The $N \times c$ indicator matrix H encodes the membership in the partition, while the $1 \times N$ vector $\pi = (\pi^{(1)}, \ldots, \pi^{(N)})$ is the stationary distribution of the random walk (1) and $\Pi = \text{diag}(\pi)$.

We look for partitions where a random walker is likely to remain trapped in the same community over the timescale t. This corresponds to large values of the diagonal elements $[R(t)]_{ii}$. Therefore, the Markov Stability of a partition H is defined as

$$r(t) = \max_H \text{trace}(R(t,H)), \tag{3}$$

and we search for partitions H that maximise $r(t)$. The 'time' t is a dimensionless quantity that measures the expansion of the random walk in the network acting as a dynamical resolution parameter, and does not directly correspond to a biophysical timespan. To emphasise this difference, we refer to t as the *Markov time* throughout. To optimise the Markov Stability $r(t)$ for any given t, we use the Louvain algorithm [41], a highly efficient greedy algorithm that finds optimised partitions with high values of $r(t)$ with no guarantees of global optimality (the problem is NP-hard), but which works well in practice.

Robustness as a measure of significant graph partitions. At any Markov time t, the above process finds an optimised partition but there is no *a priori* reason why this partition need have any significance. Indeed, it is likely that significant partitions will only be found at certain scales (e.g., at the level of chemical groups or of amino acids). Through the use of surrogate chemical randomisations [37], we have shown that Markov Stability is able to detect chemical groups, biochemical units, as well as structural features such as helical turns. At long Markov times, we look for significant partitions with the defining feature of being robust to perturbations [42]. We use two measures to quantify this robustness:

1. The length of the Markov time over which the partition is found optimal is an indication of its persistence for flow retention under small parametric changes in time. Hence we look for plateaux in the number of communities versus Markov time.

2. A widespread similarity between the partitions found by the optimisation algorithm indicates the robustness of the partition found. At each Markov time, we obtain 100 optimised partitions using 100 different initial conditions for the optimisation algorithm. We then calculate the average pairwise difference between these 100 partitions using the *variation of information* (VI) [43], an information-theoretic measure that quantifies the similarity/dissimilarity of two partitions of the same network. A low VI (or a dip in the VI) reflects greater homogeneity among the 100 optimised solutions obtained and therefore increased robustness to the optimisation.

***In silico* mutational analysis to identify significant residues.** Alanine scanning mutagenesis is the systematic experimental replacement of individual amino acids with alanine in the primary structure of a protein. We can mimic computationally the effect of this procedure by removing the graph edges corresponding to the weak interactions formed by the side chain of the chosen amino acid. The computational efficiency of our algorithm means we can evaluate the effect on the Markov Stability partitions of all individual mutations of each residue in turn. To identify the Markov times at which a mutation has a significant impact on the robustness of a partition, we compare the VI(t) graphs of the wild type versus that of the mutated networks. Gaussian Process Regression[44] is used to obtain a representative VI curve for the ensemble of mutated and wild-type VI graphs. If the VI graph of a mutated network falls outside the statistical bounds of the ensemble trajectory, it indicates that the robustness of the partition has been significantly affected (Fig. S2). We identify the point mutations that lead to significant changes in the structurally relevant graph partitions. See section S3 for a full discussion of Gaussian Process Regression.

2.4 Markov transient analysis and signal propagation

To identify special regions in the protein that are significantly connected to the active site we perform an explicit analysis of Markov transients from initial conditions localised on a particular subgraph and establish a measure of the convergence of another target subgraph towards stationarity.

Source-target transients of the random walk. Consider the evolution of a random walk described by Eq. (1). To model the propagation through the network of a perturbation occurring at a particular site, we analyse a random walk originating at that site and define an initial probability distribution $\mathbf{p}_0$ in which the probability is spread uniformly over a set of *source nodes*:

$$\mathbf{p}_0 = \Big(0 \quad \dots \quad \overbrace{\left[\tfrac{1}{N_s} \cdots \tfrac{1}{N_s}\right]}^{\mathbf{p}_0^{\text{source}}} \quad \dots \quad \overbrace{[0 \dots 0]}^{\mathbf{p}_0^{\text{target}}} \quad \dots \quad 0\Big),$$

where N_s is the number of source nodes. We then monitor signal propagation between two defined regions of the network by observing the change in probability at the *target nodes*:

$$\mathbf{p}_t = \left(\dots \left[\mathbf{p}_t^{\text{source}}\right] \dots \left[\mathbf{p}_t^{\text{target}}\right] \dots\right), \tag{4}$$

As the Markov time $t \to \infty$, the vector $\mathbf{p}_t$ converges to the stationary distribution π. Hence the speed at which the target nodes reach stationarity can be used as a measure of connectivity between the source and target nodes. See Fig. 2 for a sketch of this procedure.

Characteristic transient time $t_{1/2}$ as a measure of signal propagation between two sites. To measure the connectivity between two sites in the protein, we introduce the characteristic transient time $t_{1/2}$, a measure of the speed with which a random walk originating at one site will propagate to the target site. Given a set of source atoms, the $t_{1/2}$ associated with the target node i is the number of time steps it takes for the

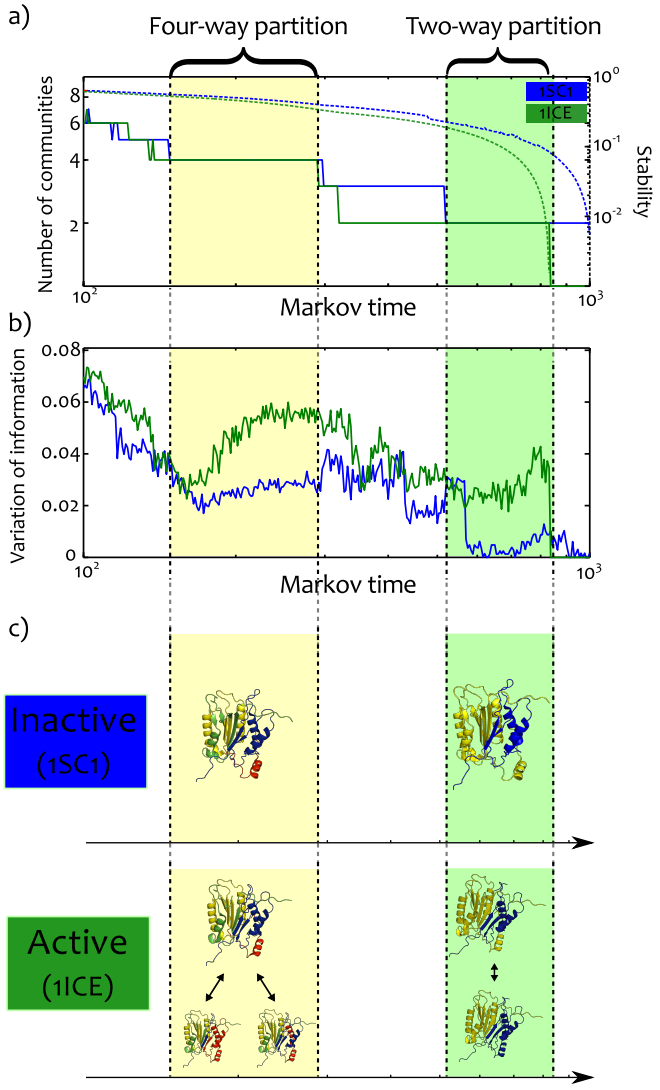

Fig. 3 Community detection across timescales for caspase-1. a) Number of communities (solid line) and Markov Stability (dashed line) between Markov times 100 and 1000 of two conformations of caspase-1: active (green) and inactive (blue). b) The variation of information of both conformations over the same time scale indicates that the 4-way and 2-way partitions of the active conformation are less robust than for the inactive conformation. c) The dominant 4-way and 2-way partitions for the two conformations. The inactive conformation has a single dominant partition, the inactive conformation flips between several different partitions.

probability at node i to reach half its stationary value:

$$t_{1/2}^{(i)} = \arg\min_t \left[p_t^{(i)} \geq \frac{\pi^{(i)}}{2} \right].$$

To measure the connectedness between two *sets* of atoms, we average the $t_{1/2}^{(i)}$ over all atoms in the target set (e.g., over the atoms of a residue or over the atoms of a group of residues)

Measuring changes in signal propagation: the $t_{1/2}$ ratios. We will be interested in measuring the changes in signal propagation from the active site to any other residue (or groups of residues) in the protein as quantified using the characteristic transient time $t_{1/2}$ defined above. We define three ratios that allow us to compare the $t_{1/2}$ under different changes in the protein:

Conformational $t_{1/2}$ ratio: measures the change in $t_{1/2}$ for each residue for a random walk originating at the active site when comparing the active and inactive conformations of the protein

$$\Delta_{\mathrm{CF}} = \frac{t_{1/2}^{\mathrm{inactive}}}{t_{1/2}^{\mathrm{active}}}. \quad (5)$$

A high ratio indicates that a residue is more closely coupled to the active site in the active conformation.

Bond-removal $t_{1/2}$ ratio: measures the contribution of individual bonds to signal propagation (between the active site and a target) by comparing the $t_{1/2}$ before and after removal of a bond in the graph

$$\Delta_{\mathrm{BR}} = \frac{t_{1/2}^{\mathrm{bond\text{-}rem}}}{t_{1/2}^{\mathrm{active}}}. \quad (6)$$

Mutational $t_{1/2}$ ratio: measures the importance of individual residues to communication between two sites by comparing the $t_{1/2}$ before and after mutation of that residue (i.e. by removal of *all* weak interactions formed by that residue)

$$\Delta_{\mathrm{MT}} = \frac{t_{1/2}^{\mathrm{mut}}}{t_{1/2}^{\mathrm{active}}}. \quad (7)$$

3 Results and discussion

3.1 Markov Stability reveals a strongly compartmental community structure in the inactive conformation of caspase-1

We have used Markov Stability to analyse the 'active' (1ICE) and 'inactive' (1SC1) structures of caspase-1. The 1CE active structure is used as it does not have unresolved residues causing a break in the backbone chain. Note that the ligand in 1ICE is not included in the graph, so as to make the comparisons with the unliganded 1SC1 consistent, i.e., the observed differences between the conformations are due to changes in graph properties induced on the protein structures and not due to the extra atoms/bonds of the ligand.

As described above, Markov Stability zooms across different levels of resolution to find robust graph communities at all scales from the atomic graph of the protein. Such communities can be thought of as groups of atoms behaving coherently over a particular timescale under a diffusive process. The full Markov Stability analysis is shown in Fig. S1. Drops in the variation of information (VI), indicative of robust partitions, are observed around Markov times 2×10^{-3} and 10^{-1} corresponding to chemical groups and amino acids, respectively, as discussed elsewhere[37]. Between Markov times 3×10^{-1} and 10, high VI and a lack of plateaux indicates an absence of significant partitions in the protein structure. The VI curve begins to fall again after Markov time 10, and around this time we see the emergence of the secondary structure with alpha-helices forming communities. These features are observed in all proteins studied with this method so far[37], and confirm that we uncover the expected structure for proteins at small and intermediate scales. Consequently, for these short and intermediate Markov time scales, there are no differences between the active and inactive conformations of caspase-1 since they share their chemical constituents.

Above Markov time 100, however, the results for the active and inactive conformations diverge and we see distinct features for each conformation of caspase-1 (Fig. 3). In particular, the inactive conformation exhibits long plateaux in the number of communities between Markov times 120-300 and 460-800, accompanied by a drop in the VI, indicating the presence of robust four-way and two-way partitions (Fig. 3). On the other hand, the active conformation presents much weaker indications of community structure, i.e., the quantitative Markov Stability curve of the inactive structure (which measures the strength of the partition) is consistently smaller and the VI is larger during the plateaux.

In particular, the four-way partition in the inactive conformation is more robust, with an average pairwise VI $\simeq 0.0254$ in the inactive structure compared to the larger 0.0435 in the active structure. In the inactive structure there are four clearly demarcated communities which comprise the four-way partition (Fig. 3c top left). In contrast, in the active conformation we identify 7 smaller sub-communities which combine in different ways to form different four-way partitions in a more flexible manner (Fig. 3c bottom left).

Furthermore, the long-lived, robust (VI $\simeq 0.0073$) two-way partition of the inactive conformation splits caspase-1 into its p10 and p20 subunits (Fig. 3c top right). In the active conformation, however, the two subunits are less well defined as

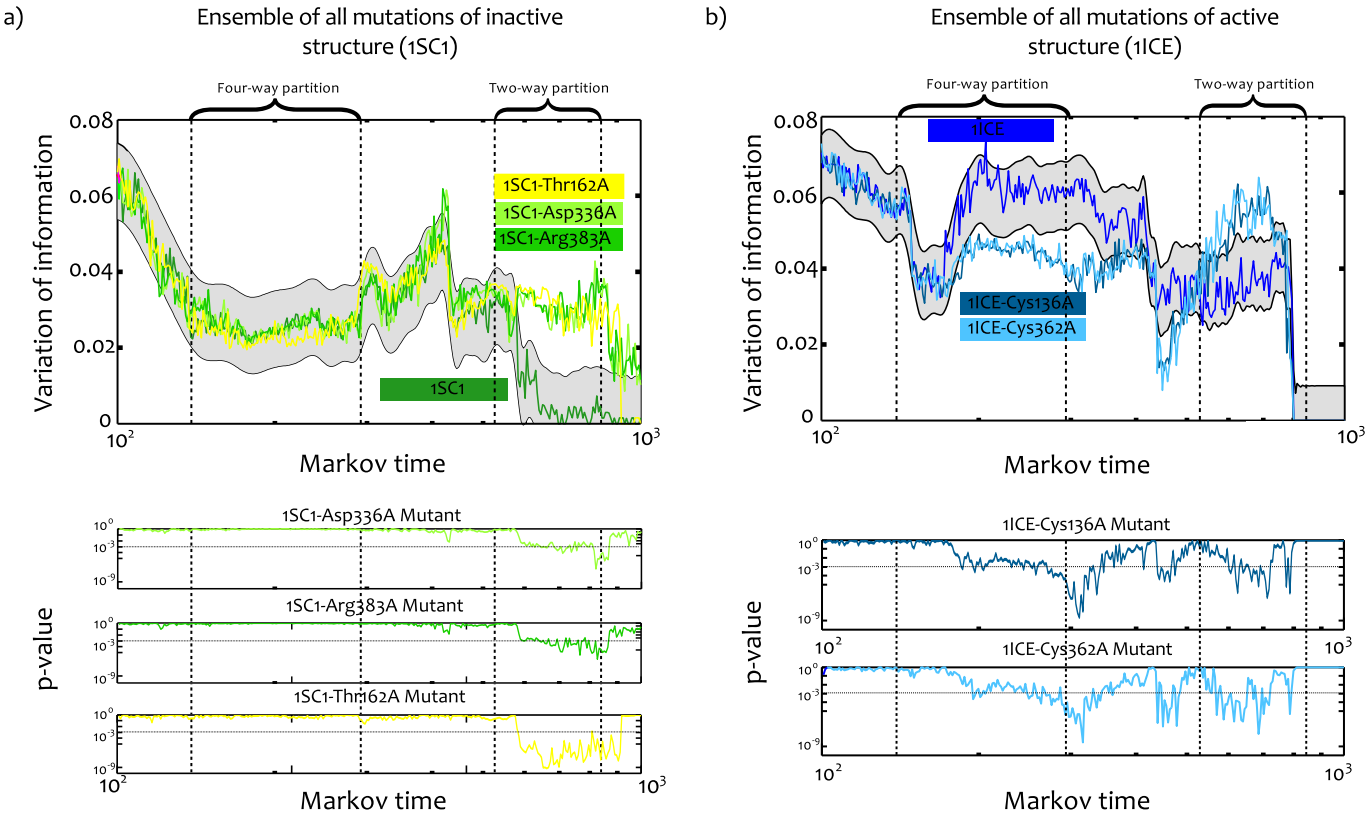

Fig. 4 Variation of information (VI) as a function of Markov time for mutated structures a) Analysis of all the mutated inactive structures: (top panel) VI as a function of Markov time with statistical bounds (95% confidence intervals) obtained with GPR applied to the ensemble of all mutations. Above Markov time 570 (corresponding to the 2-way partition), the VI for mutations of Thr162, Asp336 and Arg383 are well above the statistical bounds of the ensemble, reflecting a loss of well-defined community structure in the two-way partition similar to what is observed in the active structure; (bottom panel) The p-value quantifying the likelihood that the mutated structure does not belong in the ensemble of mutations drops sharply ($p < 0.001$) over the 2-way plateau. b) Analysis of all the mutated active structures: the VI for mutations of Cys136A and Cys362A drop below the statistical bounds between Markov times 170 and 400, representing an increased robustness of the four-way partition observed during this period (top panel). These mutations also induce changes in the robustness of the two-way partition as is also seen in the p-values shown in the bottom panel.

separate communities (VI $\simeq 0.0281$): the $\alpha - 1/2$ helix and the $\beta - 6$ strand of the p20 subunit are closely associated with the p10 subunit (Fig 3c. bottom right), indicating a stronger interaction between the p10 and p20 subunits in the active conformation.

3.2 Computational mutagenesis reveals important residues for community structure

To mimic *in silico* the process of alanine mutagenesis, we remove all edges corresponding to interactions of a given 'mutated' residue. We consider all point mutations in turn, and compute the community structure using Markov Stability for each mutated structure. We can then identify the mutations that affect significantly the robustness of the 4-way and 2-way partitions by analysing the VI of the ensemble of mutated networks using Gaussian Process Regression[44] (see Section S3). Mutations are classified as significant if the VI of the mutated structure lies 3 standard deviations outside the mean for at least one third of the relevant Markov time plateaux. Using this criterion, we find that only three residues (Thr162/Asp336/Arg383) in the inactive conformation and two residues (Cys136/Cys362) in the active conformation affect the community structure significantly at long Markov times (Fig. 4).

In the active structure, mutations of residues Cys136 and Cys362 significantly affect the four-way partition (Fig. 4b). These residues form a disulphide bond linking the p10 and p20 subunits in the active structure, which is absent in the inactive structure (Fig. S4a). Removing this bond by our computational mutation breaks a strong link between the two subunits and appears to stabilise the four-way partition (Fig. 3c),

resulting in the more compartmentalized community structure characteristic of the inactive conformation. Although experimental mutations of either residue have not been shown to have an effect on enzyme activity[34], they may affect dynamical and structural features of the protein.

The three mutations that have an effect on the large scale organisation of the inactive structure are Thr162, Asp336 and Arg383. Thr162 lies in the $\alpha - 1/2$ helix in the p20 subunit and forms hydrogen bonds with Glu223 and Thr226 (Fig. S4b). Removal of these interactions has a significant effect on the two-way partition of the inactive conformation at long Markov times (Fig. 4a): instead of a clear two-way partition into the two subunits, we find that the $\alpha - 1/2$ helix of the p20 subunit forms a community with the p10 subunit.

Asp336 and Arg383 form a salt bridge (Fig. S4c) and, similarly, the removal of this bond by mutation of either residue causes a greater association of the $\alpha - 1/2$ helix and the p10 subunit. It has been suggested that the main role of Asp336 is to stabilise the L3 loop (residues 332-346) containing residue Arg341, an important residue for substrate-recognition[31]. Our results suggest that the Asp336/Arg383 bond may play a role in stabilising the inactive conformation and that removing this salt bridge may facilitate adoption of the active conformation. Romanowksi *et al*[35] suggest that the primary function of the L4 loop, containing Arg383, may be to stabilise loop L3 in the inactive conformation. Our analysis suggests that mutation of this residue may have a global effect on the dynamics of the enzyme. In section 3.4 we find that this residue may play a role in active-site to active-site co-operativity.

3.3 Allosteric pathways uncovered by transients of random walks on the atomistic graph of the protein

The conformational $t_{1/2}$ ratio Δ_{CF} reveals the location of the allosteric site. To investigate signal propagation between the active site and the rest of the protein, we consider first the behaviour of a random walk originating at the active site. The active site is defined as the residues containing an atom within 4Å of the substrate and comprises residues 179/236-238/283-285/338-343/345/348/383[34]. We use the 2HBQ active structure for the transient analysis as this structure was used previously to characterise an allosteric hydrogen bonding network[31].

As would be expected from its weaker community structure, the random walk spreads more rapidly through the protein in the active conformation. The average characteristic transient time to any residue in the protein for a random walk originating at the active site in the active conformation is ($\bar{t}_{1/2}^{\text{active}} = 1367$), much shorter than that of the inactive conformation ($\bar{t}_{1/2}^{\text{inactive}} = 1663$). Crucially, the differences in transient times between the inactive and active conformations are not distributed homogeneously across the structure: mapping

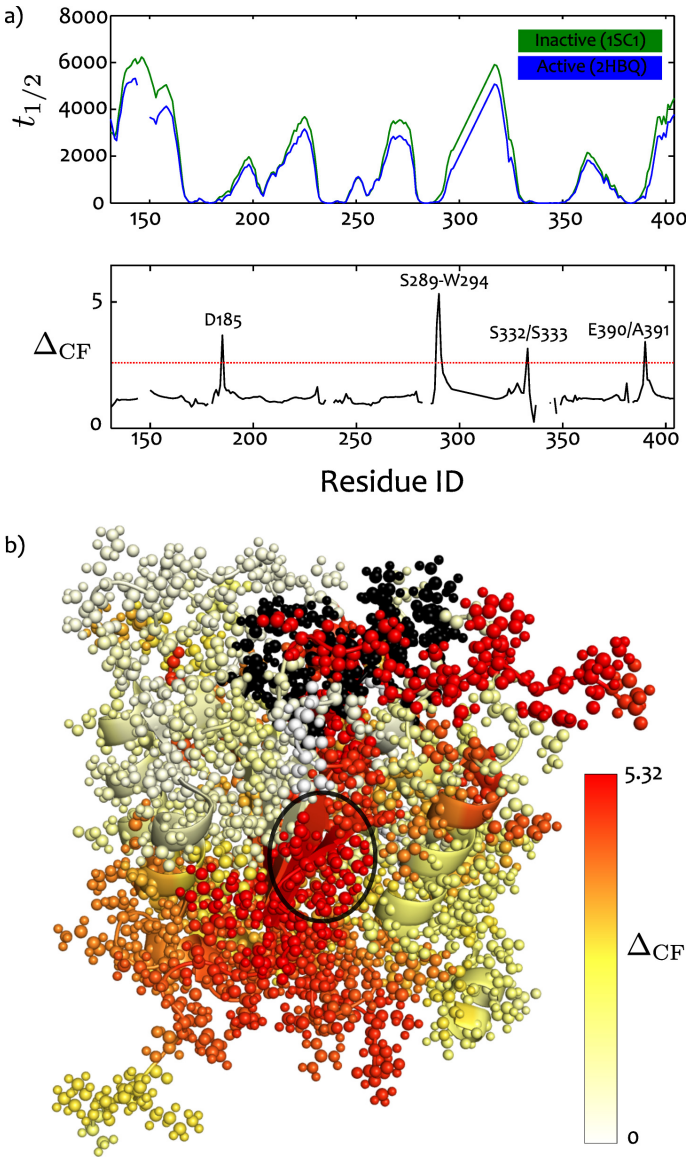

Fig. 5 Difference in signal propagation originating from the active site between active and inactive conformations. a) Characteristic transient times $t_{1/2}$ for random walks originating at the active site for all residues for the active (green) and inactive (blue) conformations (top panel), and the conformational ratio Δ_{CF} defined in Eq. (5) (bottom panel). The random walk propagates more quickly in the active conformation with $t_{1/2}^{\text{inactive}} > t_{1/2}^{\text{active}}$ consistently. The red dashed line corresponds to $\Delta_{CF} > 2$ and identifies the residues shown in Table 1. b) The conformational ratio Δ_{CF} mapped onto the protein structure: red areas show the biggest increase of $t_{1/2}$ in the inactive conformation with a 'hot spot' at the allosteric site (circled). The source atoms of the random walk (the active site) are coloured black.

Table 1 Residues with largest conformational ratio Δ_{CF} for random walks originating at the active site. Residues marked with an asterisk appear in the hydrogen bonding network identified by Datta *et al*[31].

Residue	$t_{1/2}^{\text{active}}$	$t_{1/2}^{\text{inactive}}$	Δ_{CF}
P290	50.57	268.93	5.32
S289	37.55	154.45	4.11
D185	74.50	273.75	3.67
E390*	199.60	683.73	3.43
S333*	20.09	63.36	3.15
G291	132.57	384	2.90
V292	246.25	538.25	2.19
S332*	11.18	23.82	2.13
V293	472.25	950	2.01
A391	674.96	1297.4	1.92

Table 2 Residues with largest conformational ratio Δ_{CF} for random walks originating at residue Glu390. Residues marked with an asterisk appear in the hydrogen bonding network identified by Datta *et al*[31].

Residue	$t_{1/2}^{\text{active}}$	$t_{1/2}^{\text{inactive}}$	Δ_{CF}
R286*	1	493	493
G287	36	729	20.25
D288	65	915	14.08
L258	2	25	12.5
S289	114	1159	10.17
S339*	100	930	9.3
C285	36	311	8.64
S333*	8	66	8.25
P290	193	1440	7.46
R240	244	1762	7.22
T334	18	127	7.06
G291	274	1672	6.10
W340	213	1265	5.94

the conformational ratio Δ_{CF} onto the protein structure reveals a 'hot spot' at the allosteric site (Fig. 8). In Table 1, we show the ten residues with $\Delta_{CF} > 1.9$, corresponding to the largest change between inactive and active conformations. Three of these ten residues (Ser332/Ser333/Glu390) are in the hydrogen bonding network by Datta *et al*[31]. Of these, Glu390 is notable for being located in the allosteric binding pocket at the dimer-dimer interface. Glu390 forms a salt bridge with Arg286, a bond which is known to be disrupted through allosteric inhibition. Other residues with large conformational ratios Δ_{CF} are Asp185 and residues 289-293. The large Δ_{CF} of residues 289-293, which are located in the highly dynamic C-terminus of the p20 subunit, is due to the loss of hydrogen bonds between Ser289 and two active-site residues Asp336 and Val338. Similarly, Asp185 loses a hydrogen bond with the active site residue 179.

To test the relevance of the allosteric site residues identified, we performed the transient analysis of the 'reverse' random walk originating at the residue Glu390 in the allosteric site. The random walk spreads much more quickly towards the active site in the active conformation (Fig. 6 and Video 1 in the SI), revealing the existence of a communication pathway between the allosteric and active sites which is present in the active conformation and suppressed in the inactive conformation. To identify the scaffold of this communication pathway, we compute the conformational ratio Δ_{CF} for this reverse random walk. As shown in Table 2, the largest Δ_{CF} corresponds to residue Arg286, a consequence of the formation of the salt-bridge with Glu390. Residues 285-291, adjacent to residue Arg286 and including the catalytic residue Cys285, also see significant increases. Also notable is residue Ser339 and residue Ser333, which are in the hydrogen bonding network identified by Datta *et al*.

The bond-removal $t_{1/2}$ ratio Δ_{BR} identifies bonds involved in the allosteric hydrogen-bonding network. In order to quantify further the contribution of individual bonds to this communication pathway, we calculate the bond removal ratio Δ_{BR} for a random walk originating at Glu390 with target at the active site. Table 3 shows the bonds whose removal induce the largest increases in the characteristic transient time $t_{1/2}$. The Glu390/Arg286 salt-bridge has the largest impact, reflecting the importance of this major pathway. Furthermore, five of the top six bonds identified in Table 3 are formed within residues known experimentally to have the greatest impact on catalytic activity[31] (Fig. 7a).

However, the spread of the random walk shown in Figure 6 and the high importance of bonds which do not belong to the previously identified allosteric network points to the existence of minor pathways between the two sites. In this respect, we also identify four novel interactions which cause an increase in $t_{1/2}$ of a similar magnitude: Ser236/Gln283, Arg240/Asp336, Arg341/Thr180, and Arg286/Asn337 (Fig. 7b).

Arg341 undergoes significant rearrangement between the inactive and active conformations, transferring from the surface of the protein to the substrate-binding pocket and conferring selectivity on the substrate through charge-charge interactions[34]. Interestingly, it is conserved across all human caspases. The hydrogen bond it forms with Thr180 may be important for stabilising its position in the substrate binding pocket. Thr179 and Arg341 are important substrate binding residues which provide the aspartate recognition function of caspase-1.

Asp336 may stabilise the L3 loop (residues 332-346) in the active conformation[31]. Our analysis here suggests that the bond it forms with Arg240 is important for maintaining contact with the 236-238 loop containing residue His237 and with Arg286, which is adjacent to the catalytic residue Cys285. However, both of these residues were found to have only a

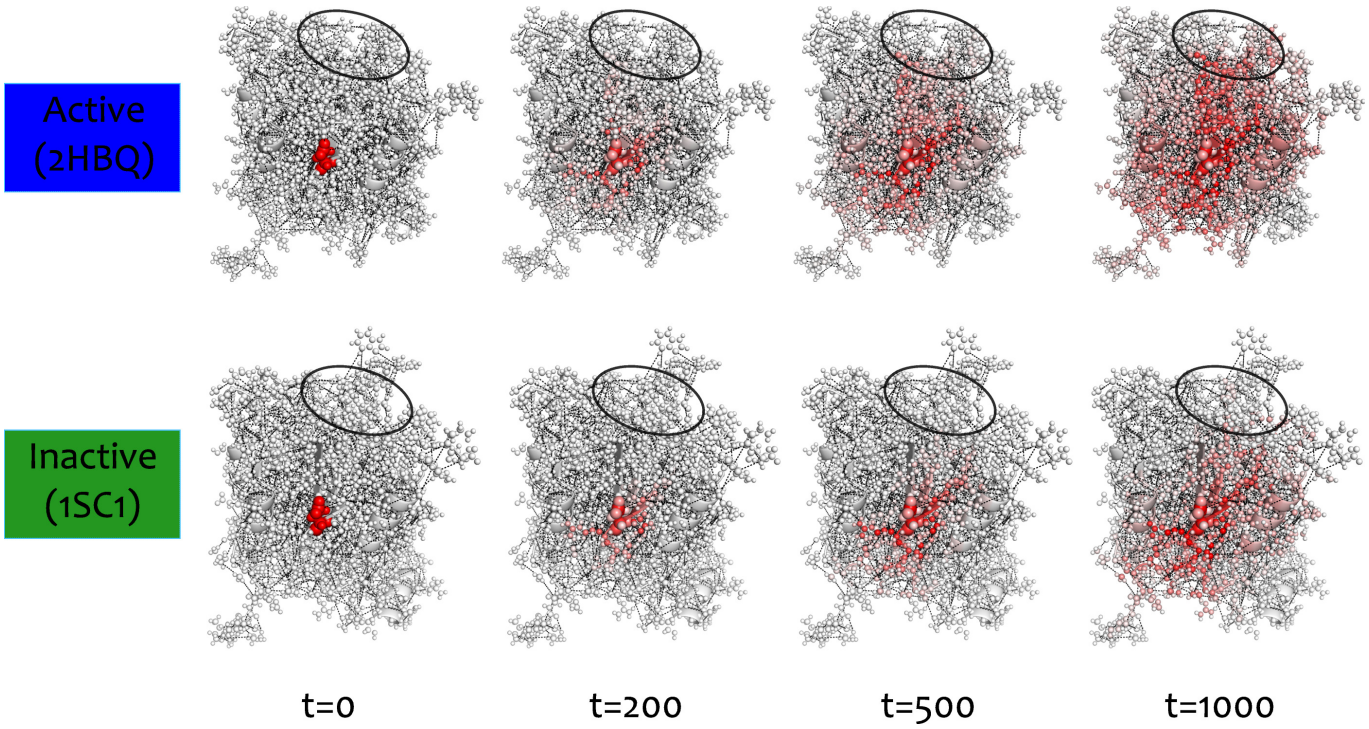

Fig. 6 Propagation of a random walk originating at Glu390. The probability of the random walk originating from residue Glu390 is shown for time steps $t = 0, 200, 500, 1000$. Atoms are coloured from white to red with increasing probability. The spread of the random walk in the active structure is broader, reflecting the weaker community structure, with increased communication with the active site (circled). On the other hand, the inactive conformation exhibits impaired communication with the active site. These images are snapshots from a video which can be viewed online.

small effect on protein function and so it would seem that these bonds are not crucial for maintaining the active conformation.

Gln283 is an active site residue which forms a hydrogen bond with a substrate sidechain[34]. Although the Ser236/Gln283 hydrogen bond is conserved between conformations, it is weaker in the active conformation. In the inactive conformation Gln283 forms a hydrogen bond with Ser347 (not preserved in the active conformation) which serves as an anchor point for the L2 loop containing the catalytic residue Cys285[35]. Thus the weakening of the Ser236/Gln283 bond could allow Gln283 to rearrange and the L2 loop to adopt its catalytically competent conformation. Our analysis suggests that the Ser236/Gln283 bond lies on a signalling pathway between the allosteric and active sites and so perturbations induced by binding of the inhibitor could affect this conformational re-arrangement. Further experimental work is required to identify whether Thr180, Gln283, or Ser236 are functionally significant.

Table 3 Bonds with largest bond removal ratio Δ_{BR} for a random walk originating at Glu390 towards the active site. Bonds marked with an asterisk are involved in the hydrogen bonding network identified by Datta *et al*[31]

Bond	Δ_{BR}
Arg286H$^{\eta 22}$: Glu390$^{\varepsilon 1}$*	1.0693
Arg286 H$^{\eta 12}$: Glu390 O$^{\varepsilon 1}$*	1.0602
Ser236 H : Gln283 O	1.0545
Arg286 H$^{\eta 11}$: Ser333 O*	1.0454
Ser332 H$^{\gamma}$: Ser339 O$^{\gamma}$*	1.0440
Ser333$^{\gamma}$: Ser339 O$^{\gamma}$*	1.0304
Arg240 H$^{\eta 11}$: Asp336 O$^{\delta 1}$	1.0294
Arg341 H$^{\eta 11}$: Thr180 O	1.0264
Arg286 H : Asn337 O	1.0256

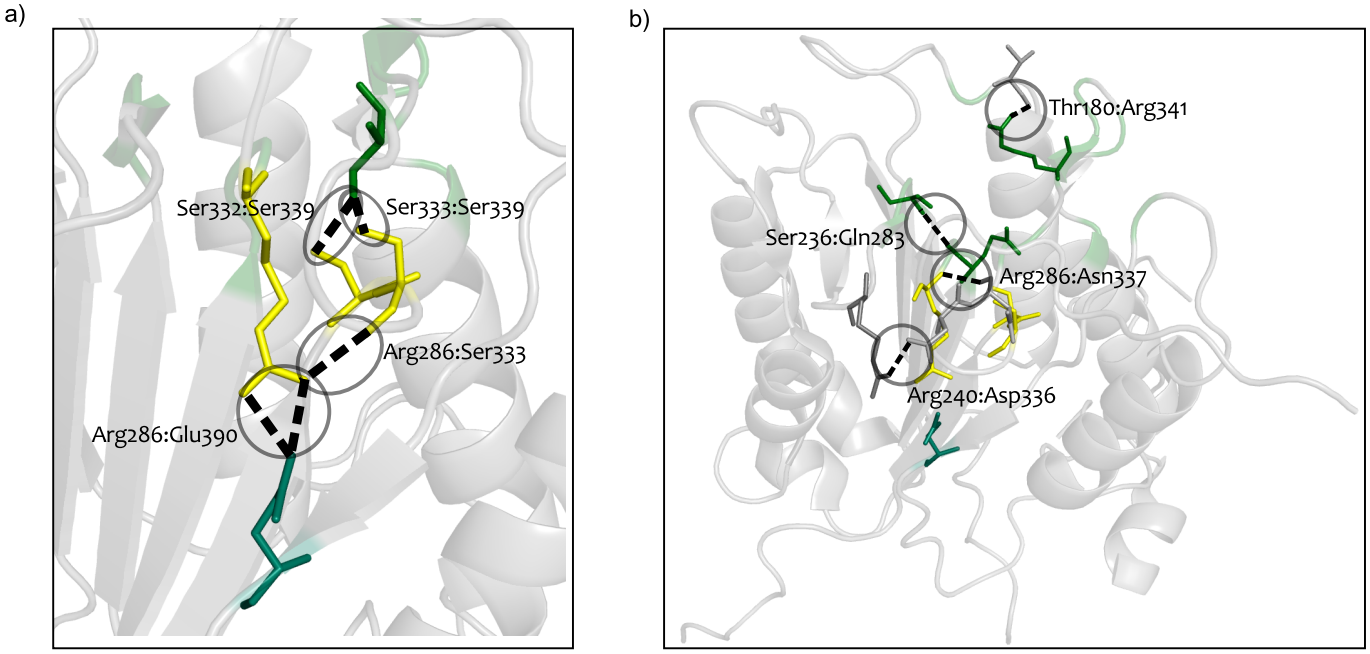

Fig. 7 Interactions identified by their high bond removal ratio Δ_{BR} as mediating communication between allosteric and active sites. In (a), we show bonds with high Δ_{BR} which are involved in the previously identified hydrogen bonding network and are known to have functional importance. Previously unidentified interactions, which we suggest help mediate allosteric communication are shown in (b): Arg341 $H^{\eta 11}$: Thr180 O, Arg341 is a crucial residue for substrate recognition; Ser236 H : Gln283 O, this bond weakens in the active form and the Gln283 sidechain rotates to bind with the substrate; Arg240 $H^{\eta 11}$: Asp336 $O^{\delta 1}$, Arg240 is adjacent to His237 which is believed to form a catalytic diad with Cys285; Arg286 H : Asn337 O, Arg286 is in the L2 loop adjacent to the catalytic residue Cys285. In both figures, the source of the random walk (residue Glu390) is coloured blue; the target residues in the active site are coloured green; and the residues involved in the known hydrogen bonding network are coloured yellow.

3.4 The mutational ratio Δ_{MT} reveals differences between active-to-allosteric-site and active-to-active-site communication in the caspase tetramer

The previous analysis has assumed no information about the location of the allosteric site, i.e., the location of the allosteric site was only used *a posteriori* to evaluate the outcomes of our algorithm. If the location of the allosteric site is known, we can integrate this information in our analysis of signal propagation. To do this, we use the mutational $t_{1/2}$ ratio (Δ_{MT}) defined in Eq. (7), which measures the impact of mutating a residue on the communication between two sites. We have performed the computational mutagenesis of all residues and calculated their $\Delta_{MT}^{\text{Act-Allost}}$ for a random walk between the active site (defined as above) and the allosteric site (defined as the residues within 3.5Å of the allosteric ligand). The top residues are shown in Table 4. Four of the key allosteric residues discussed above are identified as having large mutational ratios $\Delta_{MT}^{\text{Act-Allost}}$: E390 (1st), R286 (3rd), S339 (7th), and S332 (12th). We have also checked our results againts other independent findings. In particular, the key allosteric residues previously identified in Ref. Datta *et al.* [31] (286/332/339/390) have statistically significant higher $\Delta_{MT}^{\text{Act-Allost}}$ than other residues in the protein (Wilcoxon rank sum test $p = 6.5 \times 10^{-5}$). Furthermore, the residues with highest $\Delta_{MT}^{\text{Act-Allost}}$ correspond with those of highest functional importance, as shown by the measured impact of the mutation on the catalytic efficiency k_{cat}/K_m [31] (Fig. 8c).

In addition to allosteric inhibition, caspase-1 also exhibits strong positive cooperativity between the two active sites present in its tetramer. Recently, it has been shown that binding at one active site promotes activity at the other active-site [45], possibly due to induced dimerisation or the propagation of a conformational change. Interestingly, this cooperative behaviour is not removed by mutations implicated in allosteric inhibition [31], which suggests that cooperativity is mediated by a different mechanism. We have used our Markov transient analysis to compare the active-to-allosteric site communication versus the active-to-active site communication related to cooperativity between the two active sites of the caspase tetramer. To study the relevant residues involved in cooperative behaviour we have calculated the mutational $t_{1/2}$ ratios of all residues but now for a random walk *between the two active sites*, $\Delta_{MT}^{\text{Act-Act}}$. In Figure 8, we show that the $\Delta_{MT}^{\text{Act-Act}}$ are

Table 4 Residues with largest mutational ratio for a random walk between between the active and allosteric sites $\Delta_{MT}^{\text{Act-Allost}}$ and between the two active sites $\Delta_{MT}^{\text{Act-Act}}$ in the caspase-1 tetramer. Residues marked with an asterisk appear in the hydrogen bonding network identified by Datta *et al*[31].

Active-Allosteric		Active-Active	
Residue	$\Delta_{MT}^{\text{Act-Allost}}$	Residue	$\Delta_{MT}^{\text{Act-Act}}$
E390*	1.2590	R383	1.0845
C285	1.2275	E390*	1.0454
R286*	1.2169	R286*	1.0387
L258	1.1584	E378	1.0338
I282	1.0972	S339*	1.0326
F262	1.0890	C285	1.0288
S339*	1.0855	T389	1.0274
Q283	1.0767	R391	1.0262
I261	1.0737	L325	1.0253
I243	1.0683	N259	1.0244
L256	1.0678	T334	1.0194
S332*	1.0669	F439	1.0169

notably different to the $\Delta_{MT}^{\text{Act-Allost}}$ ratios obtained above for the active-site-to-allosteric-site random walk. For instance, the active-to-allosteric mutational ratio of E390 and R286 drops dramatically, whilst other residues increase in importance.

The residue with largest $\Delta_{MT}^{\text{Act-Act}}$ for active-site to active-site communication is Arg383. In Section 3.1, we found that the strongly compartmental community structure of the inactive conformation was weakened by mutation of Arg383. Here we find that this residue is also important for active-site to active-site communication. Residues forming dimer-dimer contacts also play a significant role in communication between the two active sites. In addition to Arg383, the residue pairs E378/L325 and T389/R391, which form weak interactions between the two dimers in the active conformation, also have high $\Delta_{MT}^{\text{Act-Act}}$. This suggests that such connections are important for transmitting signals between the two active sites and these residues could therefore play a role in transmitting binding-induced conformational changes.

4 Conclusions

In this paper we have used two methods that exploit the transients of a Markov process diffusing on an atomistic biophysical graph derived from protein structures in order to study allosteric communication pathways. Firstly, the Markov Stability community detection method identifies differences in the multiscale structural organisation of the active and inactive structures of caspase-1. In contrast to its inactive counterpart, the active conformation exhibits a fluid, weakly compartmentalised community structure with less robust partitions at large scales. This suggests that perturbations propagate over

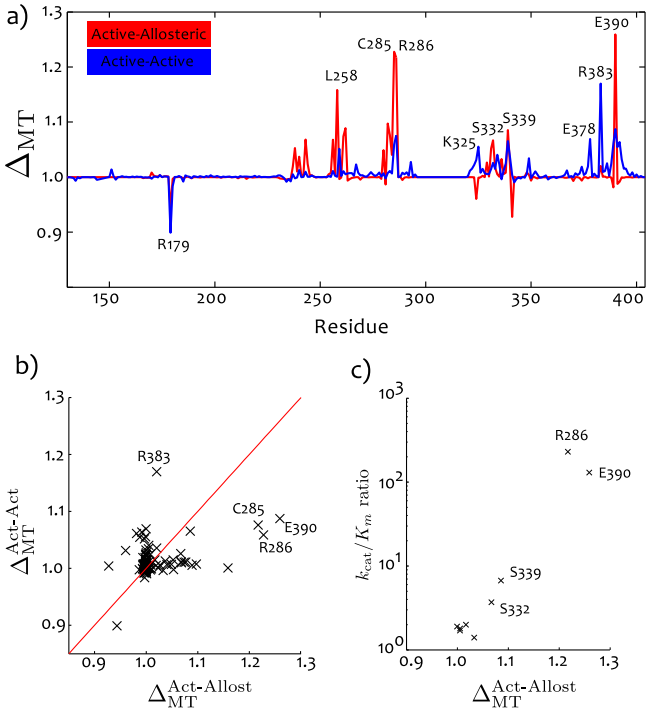

Fig. 8 Comparing the mutational $t_{1/2}$ ratios between allosteric-to-active-site and active-to-active-site communication. (a) The mutational $t_{1/2}$ ratio for each residue for random walks between the active-site and allosteric-site (red) and between the two active-sites (blue) show differences. (b) The active-site/active-site ratio $\Delta_{MT}^{\text{Act-Act}}$ plotted against the active-site/allosteric-site ratio $\Delta_{MT}^{\text{Act-Allost}}$: the $\Delta_{MT}^{\text{Act-Act}}$ of residues implicated in allosteric inhibition (E390/R286) is much lower than their active-site to allosteric-site $t_{1/2}$ ratio, which suggests that these residues are less important for co-operative behaviour. (c) $\Delta_{MT}^{\text{Act-Allost}}$ plotted against the experimental k_{cat}/K_m ratio (for residues with available experimental data[31]): residues with large $\Delta_{MT}^{\text{Act-Allost}}$ correspond to those with greatest functional significance.

larger distances in the active conformation, due to an increase in long-range communication pathways. Computational mutational analysis identifies the bond between residues Arg383 and Asp336 as crucial for maintaining this modular organisation.

Secondly, the analysis of transients of random walks originating in the active site suggests the existence of long-range communication pathways towards the allosteric site, which show distinctive characteristics in the active and inactive conformations. We have introduced three related quantitative criteria, the conformational, mutational, and bond-removal $t_{1/2}$ ratios, which allow us to identify the relevant bonds and residues on these pathways by comparing the changes in their time-dependent participation in the transient. Using the conformational $t_{1/2}$ ratio allows us to detect a hot spot at the allosteric site. Many of the residues and bonds identified with

these criteria correspond to mutations of known functional significance. We have also found several novel interactions which may be involved in alternative allosteric pathways. Finally, we have used a full computational point mutagenesis to compare the allosteric-to-active site communication versus the active-to-active site communication in the caspase dimer and shown that the relevant residues for allosteric pathways are distinct from those that play a role in the communication between active sites related to cooperativity. This agrees with experimental findings that mutation of allosteric network residues does not affect cooperativity.

Our method is a computationally efficient method to study the many parallel communication pathways in biomolecules in a probabilistic setting. Measures which go beyond binary comparisons ('present' or 'absent')[31] of bonds in the active/inactive structures provide more information about the weak interactions which mediate allosteric signals. Our transient random walk approach allows such an analysis and uncovers bonds lying on multiple pathways between the active and allosteric sites. Our Markov transient analysis provides a unified understanding which brings together structural community detection and random walk pathway identification, and therefore offers a robust way to look for candidate residues with important structural and functional roles in proteins. In particular, this method can be used to identify residues lying on communication pathways between an effector site (be it allosteric or another active site) and the active site.

5 Acknowledgements

This work was funded by an EPSRC Centre for Doctoral Training Studentship from the Institute of Chemical Biology (Imperial College London) awarded to BA.

References

1 Q. Cui and M. Karplus, *Protein Science*, 2008, **17**, 1295–1307.
2 D. Koshland Jr, G. Nemethy and D. Filmer, *Biochemistry*, 1966, **5**, 365–385.
3 J. Monod, J. Wyman and J. Changeux, *Journal of molecular biology*, 1965, **12**, 88–118.
4 J. Swain and L. Gierasch, *Current Opinion in Structural Biology*, 2006, **16**, 102–108.
5 K. Gunasekaran, B. Ma and R. Nussinov, *PROTEINS: Structure, Function and Bioinformatics*, 2004, **57**, 433–443.
6 A. Cooper and D. Dryden, *European Biophysics Journal*, 1984, **11**, 103–109.
7 N. Popovych, S. Sun, R. Ebright and C. Kalodimos, *Nature structural & molecular biology*, 2006, **13**, 831–838.
8 K. Henzler-Wildman and D. Kern, *Nature*, 2007, **450**, 964–972.
9 P. Zhuravlev, G. Papoian *et al.*, *Quarterly reviews of biophysics*, 2010, **43**, 295–332.
10 A. del Sol, C. Tsai, B. Ma and R. Nussinov, *Structure*, 2009, **17**, 1042–1050.
11 S. Lockless and R. Ranganathan, *Science*, 1999, **286**, 295–299.
12 E. Fuentes, C. Der and A. Lee, *Journal of molecular biology*, 2004, **335**, 1105–1115.
13 E. Fuentes, S. Gilmore, R. Mauldin and A. Lee, *Journal of molecular biology*, 2006, **364**, 337–351.
14 Y. Kong and M. Karplus, *Proteins: Structure, Function, and Bioinformatics*, 2009, **74**, 145–154.
15 N. Ota and D. Agard, *Journal of molecular biology*, 2005, **351**, 345–354.
16 V. Hilser, B. G.-M. E., T. Oas, G. Kapp and S. Whitten, *Chemical reviews*, 2006, **106**, 1545–1558.
17 H. Pan, J. Lee and V. Hilser, *Proceedings of the National Academy of Sciences*, 2000, **97**, 12020.
18 A. Del Sol, H. Fujihashi, D. Amoros and R. Nussinov, *Molecular systems biology*, 2006, **2**, n/a.
19 A. Del Sol, M. Araúzo-Bravo, D. Amoros, R. Nussinov *et al.*, *Genome Biol*, 2007, **8**, R92.
20 A. Sethi, J. Eargle, A. Black and Z. Luthey-Schulten, *Proceedings of the National Academy of Sciences*, 2009, **106**, 6620.
21 A. Ghosh and S. Vishveshwara, *Proceedings of the National Academy of Sciences*, 2007, **104**, 15711.
22 C. Chennubhotla and I. Bahar, *PLoS computational biology*, 2007, **3**, e172.
23 K. Park and D. Kim, *BMC bioinformatics*, 2011, **12**, S23.
24 H. Lu and J. Liang, *PLoS Computational Biology*, 2009, **5**, e1000526.
25 C. Chennubhotla and I. Bahar, *Molecular systems biology*, 2006, **2**, n/a.
26 J. Delvenne, S. Yaliraki and M. Barahona, *Proceedings of the National Academy of Sciences*, 2010, **107**, 12755–12760.
27 M. E. Newman, *Proceedings of the National Academy of Sciences*, 2006, **103**, 8577–8582.
28 S. Fortunato and M. Barthelemy, *Proceedings of the National Academy of Sciences*, 2007, **104**, 36–41.
29 M. T. Schaub, J.-C. Delvenne, S. N. Yaliraki and M. Barahona, *PloS one*, 2012, **7**, e32210.
30 J. Scheer, M. Romanowski and J. Wells, *Proceedings of the National Academy of Sciences*, 2006, **103**, 7595–7600.
31 D. Datta, J. Scheer, M. Romanowski and J. Wells, *Journal of molecular biology*, 2008, **381**, 1157–1167.
32 P. Sleath, R. Hendrickson, S. Kronheim, C. March and R. Black, *Journal of Biological Chemistry*, 1990, **265**, 14526–14528.
33 J. Li and J. Yuan, *Oncogene*, 2008, **27**, 6194–6206.
34 K. Wilson, J. Black, J. Thomson, E. Kim, J. Griffith, M. Navia, M. Murcko, S. Chambers, R. Aldape, S. Raybuck *et al.*, *Nature*, 1994, **370**, 270–275.
35 M. Romanowski, J. Scheer, T. O'Brien and R. McDowell, *Structure*, 2004, **12**, 1361–1371.
36 H. Berman, J. Westbrook, Z. Feng, G. Gilliland, T. Bhat, H. Weissig, I. Shindyalov and P. Bourne, *Nucleic acids research*, 2000, **28**, 235–242.
37 A. Delmotte, E. Tate, S. Yaliraki and M. Barahona, *Physical Biology*, 2011, **8**, 055010.
38 J. M. Word, S. C. Lovell, J. S. Richardson and D. C. Richardson, *Journal of molecular biology*, 1999, **285**, 1735–1747.
39 D. Jacobs, A. Rader, L. Kuhn and M. Thorpe, *Proteins: Structure, Function, and Bioinformatics*, 2001, **44**, 150–165.
40 S. Mayo, B. Olafson and W. Goddard, *Journal of Physical Chemistry*, 1990, **94**, 8897–8909.
41 V. Blondel, J. Guillaume, R. Lambiotte and E. Lefebvre, *Journal of Statistical Mechanics: Theory and Experiment*, 2008, **2008**, P10008.
42 B. Karrer, E. Levina and M. Newman, *Physical Review E*, 2008, **77**, 046119.
43 M. Meilă, *Journal of Multivariate Analysis*, 2007, **98**, 873–895.
44 C. E. Rasmussen, *Gaussian processes for machine learning*, MIT Press, 2006.
45 D. Datta, C. L. McClendon, M. P. Jacobson and J. A. Wells, *Journal of*

Biological Chemistry, 2013, **288**, 9971–9981.